

PATENT

What is claimed is:

1 1. A switch for a packet switched or frame switched network, said switch having conventional
2 routing circuitry, said network having client devices and storage devices and a storage manager
3 that decides to redirect read and/or write transactions on any basis so as to virtualize said
4 storage devices, said switch comprising:

5 one or more ports which are structured to operate to receive one or more redirection
6 commands from said storage manager or another switch in said network coupled to said
7 storage manager containing old address data and new address data, said ports containing
8 circuitry to store said old address data and new address data in one or more look up tables,
9 said switch containing circuitry functioning to compare packet or frame type information
10 included in a header of at least some packets or frames arriving at a port coupled to a client
11 device or a storage device to determine if the packet or frame is of a type which is to be
12 relabelled and redirected, and, if the packet or frame is of a type to be relabelled and
13 redirected, comparing at least some of the old address data therein to the old address data
14 entries for one or more redirection commands stored in one or more of said look up tables,
15 and if there is a match to any of the old address data of said one or more redirection
16 commands, relabelling said packet or frame by substituting the corresponding new address
17 data from said matching redirection command(s) for said old address data in said packet or
18 frame header, said new address data being such as to cause said relabelled packet or frame
19 to be routed so as to bypass said storage manager but to look like it came from said storage
20 manager, and then passing said relabelled packet or frame to said conventional routing
21 circuitry for forwarding to the destination identified in said new address data, and wherein the
22 sequence in which the type of incoming packet is determined and the old address data
23 lookup is performed is not critical.

1 2. The apparatus of claim 1 wherein each said port contains its own redirection circuitry and
2 its own look up table for storing redirection commands, said redirection circuitry in every port being
3 structured to receive first and second redirection commands for every write transaction to be
4 redirected and to use configuration data defining which switch ports are coupled to which
5 devices, said configuration data being stored in said switch or received in a message from said
6 storage manager and said configuration data being used to forward said first redirection
7 command to a first port of said switch coupled to the client device which originated said write
8 request for storage in a look up table, and to forward the second redirection command to a
9 second port coupled to the storage device which is to store the data to be written for storage in a
10 look up table, said redirection circuitry being structured to receive a redirection command for every

PATENT

1 1 read transaction to be redirected and to use configuration data to forward said redirection
1 2 command to a third port coupled to the storage device where the requested data is stored for
1 3 storage in a look up table.

1 3. The apparatus of claim 1 wherein each said port contains its own redirection circuitry and
2 its own look up table for storing redirection commands, said redirection circuitry in every port being
3 structured to receive first and second redirection commands for every write transaction to be
4 redirected and to forward said first and second redirection commands to all ports of said switch for
5 storage in the look up tables of said ports, said redirection circuitry also being structured to
6 receive a redirection command for every read transaction to be redirected and to forward said
7 redirection command to all ports of said switch for storage in a look up table in said third port.

1 4. The apparatus of claim 1 wherein each said port contains its own redirection circuitry, said
2 redirection circuitry in every port being structured to receive first and second redirection
3 commands for every write transaction to be redirected and to forward said first and second
4 redirection commands to a shared look up table in said switch to which all redirection circuits in all
5 ports have access, said redirection circuitry also being structured to receive a redirection
6 command for every read transaction to be redirected and to forward said redirection command for
7 storage in said shared look up table.

1 5. The apparatus of claim 1 wherein said switch contains shared redirection circuitry that
2 performs the redirection process and a shared look up table for storing redirection command used
3 by said shared redirection circuitry to perform said redirection process, and wherein each said port
4 contains circuitry to store redirection commands in said shared look up table by any method, and
5 to forward frames or packets to said shared redirection circuitry for analysis and relabelling if said
6 frames or packets are to be redirected.

1 6. A virtualizing storage manager for a packet switched or frame switched network, said
2 network having client devices and storage devices and a switch having any structure capable of
3 redirecting read and/or write transactions based upon commands from said storage manager so
4 as to virtualize said storage devices, said storage manager comprising:

5 conventional storage manager circuitry to carry out conventional processing of read
6 and/or write transactions that are not redirected; and

7 a computer programmed to or logic circuitry structured to decide whether to redirect
8 packets or frames of a read or write transaction based upon any redirection criterion, and
9 programmed or structured to generate and send at least one said redirection command to a

PATENT

switch in said network for every read or write transaction to be redirected, and programmed or structured to send a corresponding read or write request to the appropriate storage device for every received read or write request from a client device to start the transaction, and programmed or structured to send at least one purge command to a switch in said network which stored said redirection commands in one or more look up tables, said at least one purge command corresponding to one or more completed read or write transactions that have been redirected by said switch to cause said switch to purge from said one or more look up tables the redirection command or commands pertaining to said one or more read or write transaction that has been completed.

7. The apparatus of claim 6 wherein said storage manager makes a decision based upon programmable redirection criteria.

8. The apparatus of claim 6 wherein said storage manager makes a decision whether or not to redirect data and transfer ready packets or frames of a particular read or write request based upon redirection criteria that is that every packet or frame of a data or transfer ready type which is part of a read or write transaction is to be redirected.

9. The apparatus of claim 6 wherein said storage manager makes a decision whether or not to redirect data and transfer ready packets or frames of a particular read or write request based upon redirection criteria that is that every packet or frame of a data or transfer ready type which is part of a read or write transaction is to be redirected when the number of outstanding read and/or write requests is greater than or equal to a threshold.

10. The apparatus of claim 9 wherein said threshold is programmable.

11. The apparatus of claim 6 wherein said storage manager makes a decision whether or not to redirect data and transfer ready packets or frames of a particular read or write request based upon redirection criteria that is that every packet or frame of a data or transfer ready type which is part of a read or write transaction is to be redirected when the number of outstanding read and/or write requests is such that a measured latency between receipt of a read or write request and completion of the corresponding read or write transaction is greater than a threshold latency.

12. The apparatus of claim 6 wherein said storage manager makes a decision whether or not to redirect data and transfer ready packets or frames of a particular read or write request based upon redirection criteria that is that every packet or frame of a data or transfer ready type which is

PATENT

4 part of a read or write transaction is to be redirected when said storage manager decides based
5 upon any type criteria or any type measurement that it has become a bottleneck.

1 13. The apparatus of claim 6 wherein said storage manager makes a decision whether or not
2 to redirect data and transfer ready packets or frames of a particular read or write request based
3 upon redirection criteria that is that every packet or frame of a data or transfer ready type which is
4 part of a read or write transaction is to be redirected when the count of a counter which was
5 started when a read or write request from a client device was received and stopped when the
6 read or write transaction corresponding to said request has been completed exceeds a
7 predetermined threshold.

1 14. The apparatus of claim 13 wherein said predetermined threshold is a number read from a
2 look up table which stores threshold values based upon the size of the read or write transaction,
3 and wherein said means for deciding performs the following steps:

4 (1) analyzing each read or write request to determine the size of the transaction and
5 retrieves the appropriate latency threshold value from said look up table based upon the size
6 of the read or write transaction;

7 (2) starting a counter when said read or write request was received;

8 (3) perform conventional read or write transaction processing including sending a read or
9 write command to the appropriate storage device and including forwarding data and transfer
10 ready frames to the client or storage device as appropriate to the particular type of
11 transaction being performed;

12 (4) stop said counter when said conventional read or write transaction was completed;

13 (5) comparing the counter value when the transaction was completed to the threshold
14 number looked up from said look up table as a measure of the time it took to complete the
15 conventional read or write transaction; and

16 (6) if the time taken to complete the conventional read or write transaction exceeded the
17 threshold, generating one or more redirection commands for subsequent read or write
18 requests and sending them to the switch;

19 (7) performing any method of stopping redirection after some time passes and resuming
20 the latency measuring process of steps (1) through (6) for subsequent read or write requests
21 until a latency threshold value is again exceeded, and then repeating steps (6) and (7).

1 15. The apparatus of claim 14 wherein step (7) comprises continuing redirection for a fixed or
2 programmable amount of time, and then ceasing redirection and resuming performance of steps

PATENT

3 (1) through (6) for subsequent read or write requests until a latency threshold value is again
4 exceeded, and then repeating steps (6) and (7).

1 16. The apparatus of claim 14 wherein step (7) comprises continuing redirection for all
2 subsequent read or write requests except every Xth read or write request where X is any integer
3 and carrying out conventional read or write transaction processing as appropriate for every Xth
4 read or write request and measuring the latency of said Xth request by performance of steps (1)
5 through (6), and if the latency threshold has not been exceeded by said Xth request, resuming
6 performance of steps (1) through (6) for subsequent read or write requests until a latency
7 threshold value is again exceeded, and then repeating steps (6) and (7).

1 17. The apparatus of claim 6 wherein said storage manager makes a decision whether or not
2 to redirect data and transfer ready packets or frames of a particular read or write request based
3 upon redirection criteria that is that every packet or frame of a data or transfer ready type which is
4 part of a read or write transaction that exceeds a certain size as determined from the original read
5 or write request is to be redirected.

1 18. The apparatus of claim 6 wherein said storage manager includes a cache memory and a
2 cache algorithm that manages said cache memory, and wherein said storage manager makes a
3 decision whether or not to redirect data and transfer ready packets or frames of a particular read
4 or write request based upon whether the requested data is already stored in said cache memory
5 in said storage manager, or makes a decision whether or not to redirect a read transaction based
6 upon whether said storage manager cache algorithm wants to copy the data of the read request
7 into said cache memory in said storage manager.

1 19. The apparatus of claim 1 wherein each port of said switch is structured to receive
2 redirection commands either from an external source or from a source internal to the switch such
3 as the routing circuitry and store all said redirection commands in a look up table maintained by
4 said port and use said look up table to do all necessary relabelling and redirection operations.

1 20. A switch for a packet switched or frame switched network, said switch having
2 conventional routing circuitry, said network having client devices and storage devices and a
3 storage manager coupled to said switch, said switch comprising:
4 one or more ports which are structured to operate to receive redirection commands from a
5 storage manager containing old address data and new address data and to store the old

PATENT

6 address data and new address data contained therein in a shared look up table, but
7 otherwise to function as conventional switch ports;

8 shared redirection interface circuitry coupled to each of said ports and said shared look
9 up table and coupled to said conventional routing circuitry, said shared redirection interface
10 circuitry structured to compare packet or frame type information included in a header of each
11 packet or frame arriving from a port coupled to a client device or a storage device to
12 determine if the packet or frame is of a type which is to be relabelled and redirected, and, if
13 the packet or frame is of a type to be relabelled and redirected, then comparing the old
14 address data therein to the old address data in said shared look up table, and if there is a
15 match to any of the old address data entries in said look up table, substituting the
16 corresponding new address data from said look up table, and forwarding said packet or frame
17 to said conventional routing circuitry for routing.

1 21. A packet switched or frame switched network, comprising:

2 one or more client devices;

3 one or more storage devices;

4 a storage manager having redirection circuitry functioning to receive read or write
5 transaction requests from said one or more client devices and redirect at least some of said
6 read and write transaction requests issued by said client devices by issuing redirection
7 commands which contain old address data and new address data which will cause relabelling
8 of data and transfer ready frames so that they are routed between said client device which
9 issued said request and the storage device which is to be read from or written to so as to
10 bypass said storage manager, said redirection circuitry also for sending a corresponding read
11 or write request to the appropriate storage device for every received read or write request
12 from a client device to start the transaction, said redirection circuitry also functioning to
13 receive status packets or frames indicating particular read or write transactions that have
14 been redirected have been completed and for issuing purge commands which function to
15 cause purging of old address and new address data of redirected read or write transactions
16 that have been completed;

17 one or more switches, each having conventional routing circuitry and each further
18 comprising:

19 a shared lookup table;

20 a shared redirection interface circuit;

21 one or more ports which are structured to operate to receive redirection

22 commands from said storage manager or another switch and to store the old address
23 data and new address data contained therein in said shared look up table, and

PATENT

2 4 structured to receive said purge commands and to use the information therein to purge
2 5 obsolete old address data and new address data from said shared look up table for read
2 6 or write transactions that have been completed, but otherwise to function as conventional
2 7 switch ports;

2 8 and wherein said shared redirection interface circuitry coupled to each of said
2 9 ports and said shared look up table and coupled to said conventional routing circuitry,
3 0 said shared redirection interface circuitry structured to compare packet or frame type
3 1 information included in a header of each packet or frame arriving from a port coupled to a
3 2 client device or a storage device to determine if the packet or frame is of a type which is
3 3 to be relabelled and redirected, and, if the packet or frame is of a type to be relabelled
3 4 and redirected, then comparing the old address data therein to the old address data in
3 5 said shared look up table, and if there is a match to any of the old address data entries in
3 6 said look up table, substituting the corresponding new address data from said look up
3 7 table, and forwarding said packet or frame to said conventional routing circuitry for
3 8 routing; and

3 9 a data path coupling each said client devices, storage devices and said storage manager
4 0 to at least one of said one or more switches and coupling said one or more switches
4 1 together.

1 22. A packet switched or frame switched network, comprising:

2 one or more client devices;
3 one or more storage devices;

4 a storage manager having redirection circuitry functioning to receive read or write
5 requests from said one or more client devices and redirect at least some of said read and
6 write transaction requests issued by said client devices by issuing redirection commands
7 which contain old address data and new address data which will cause relabelling of data
8 and transfer ready frames so that they are routed between said client device which issued
9 said request and the storage device which is to be read from or written to so as to bypass
1 0 said storage manager, said redirection circuitry also for sending a corresponding read or write
1 1 request to the appropriate storage device for every received read or write request from a
1 2 client device to start the transaction, said redirection circuitry also functioning to receive status
1 3 packets or frames indicating particular read or write transactions that have been redirected
1 4 have been completed and for issuing purge commands which function to cause purging of
1 5 old address and new address data of redirected read or write transactions that have been
1 6 completed;

PATENT

17 one or more switches each having conventional routing circuitry and each further
18 comprising one or more ports
19 and wherein at least the ports coupled to said storage manager or another switch
20 being structured to operate to receive one or more redirection commands from said
21 storage manager or another switch in said network coupled to said storage manager
22 containing old address data and new address data and store said old address data and
23 new address data in one or more look up tables and to receive purge commands from
24 said storage manager or another switch in said network and use the information therein
25 to purge old and new address data pertaining to redirected read or write transactions that
26 have been completed from said one or more look up tables,
27 each said switch containing redirection circuitry functioning to compare packet or
28 frame type information included in a header of at least some packets or frames arriving at
29 a port coupled to a client device or a storage device to determine if the packet or frame is
30 of a type which is to be relabelled and redirected, and, if the packet or frame is of a type
31 to be relabelled and redirected, then comparing the old address data therein to the old
32 address data entries in said look up table, and if there is a match to any of the old
33 address data entries in said one or more look up tables, substituting the corresponding
34 new address data from said look up table for said old address data in said packet or
35 frame header and passing said packet or frame to said conventional routing circuitry for
36 forwarding to the destination identified in said new address data, at least the port; and
37 a data path coupling each said client devices, storage devices and said storage manager
38 to at least one of said one or more switches and coupling said one or more switches
39 together.

1 23. The apparatus of claim 21 or 22 wherein said redirection circuitry in said storage
2 manager includes means to decide which read and/or write transactions to redirect based upon
3 predetermined redirection criteria which could be a single criteria or any one of a plurality of
4 different criteria, said criteria being either fixed or programmable, and which particular criteria is
5 used can be fixed or configurable.

1 24. The apparatus of claim 22 wherein said one or more look up tables comprise a look up
2 table in each port and wherein said redirection circuitry comprises a redirection circuit in every
3 port.

1 25. The apparatus of claim 22 wherein said one or more look up tables comprise a look up
2 table in each port and wherein said redirection circuitry comprises a redirection circuit in every

PATENT

3 port, and wherein said redirection circuitry in said storage manager includes means to decide
4 which read and/or write transactions to redirect based upon predetermined redirection criteria
5 which could be a single criteria or any one of a plurality of different criteria, said criteria being
6 either fixed or programmable, and which particular criteria is used can be fixed or configurable.

1 26. A switch for a packet or frame switched network including one or more client devices and
2 one or more storage devices and a storage manager server, said switch having conventional
3 routing circuitry and means for receiving and storing redirection commands from said storage
4 manager server and for comparing address data of predetermined types of frames or packets
5 arriving at said port to stored redirection data from said redirection commands and for
6 readdressing any frames or packets of said predetermined types which match any stored
7 redirection command and for receiving purge commands from said storage manager server and
8 purging old address data and new address data identified in said purge commands of redirection
9 commands issued for transactions that have been completed.

1 27. A storage manager for a packet or frame switched network including one or more client
2 devices and one or more storage devices and a switch, comprising:

3 means receiving read or write transaction requests from said one or more client devices
4 and mapping the data designated in said read or write request to a storage device, and for
5 sending corresponding read or write requests to said storage device to start the transaction;
6 and

7 means for using redirection criteria to decide which read and write requests from said
8 client devices to redirect and for generating and sending redirection commands to said switch
9 for the read and write requests to be redirected, said redirection commands containing old
10 address data and new address data which will cause data and transfer ready frames of read
11 and write transactions to be redirected to be relabelled so that they are routed by said switch
12 to bypass said storage manager but look like they came from said storage manager; and

13 means for receiving status frames or packets and for issuing purge commands to said
14 switch when said status frames or packets indicate a read or write transaction which has been
15 redirected has been completed, said purge commands for causing said switch to purge from
16 one or more look up tables in said switch old address data and new address data identified in
17 said purge commands of redirection commands previously issued for transactions that have
18 been completed.

1 28. A packet or frame switched network including one or more client devices and one or
2 more storage devices each of which is coupled to a packet or frame switch, comprising:

PATENT

3 a switch having conventional routing circuitry therein, said switch comprising:
4 one or more ports in said switch, said switch including redirection circuitry
5 structured to operate to receive redirection commands from a storage manager
6 and to store the redirection address data containing old address data and new
7 address data contained therein in a look up table, said redirection circuitry also
8 structured to receive incoming packets or frames and determine if they are data or
9 transfer ready type frames or packets and, if not to transfer said packet or frame
10 to said conventional routing circuitry but if so, to compare the address data
11 therein to the old address data in said look up table, and, if there is a match to
12 any of the old address data in said look up table, functioning to substitute the
13 new address data from said look up table from the entry which includes old
14 address data that matched the address data in the incoming packet or frame and
15 pass the packet or frame to said conventional routing circuitry; and

16 a storage manager comprising a one or more port adapter circuits, a memory and a
17 microprocessor coupled to said memory and programmed to receive read or write requests
18 from said one or more client devices and map the data designated in said read or write
19 request to a particular storage device and to send corresponding read or write requests to
20 said storage device to start the transaction, and programmed to determine from the frame or
21 packet header information of said read or write request the specific client device from which
22 the request came and how large the transaction is, and, if the transaction is larger than a
23 redirection size criteria, for generating and sending one or more redirection commands to said
24 switch for storage in at least one look up table maintained therein, said redirection command
25 containing old address data and new address data that will cause said switch to receive data
26 and transfer ready frames for transactions to be redirected to relabel said packet or frame
27 with new address data so that said data or transfer ready frames will be routed by said switch
28 so as to bypass said storage manager but look like they came from said storage manager,
29 and for performing conventional processing on any read or write request that is not to be
30 redirected, and programmed or structured to monitor status frames to determine when a
31 redirected transaction has been completed and sending purge commands to said switch to
32 purge the old and new address data from said look up table for read and write transactions
33 that have been completed; and

34 a data path coupling said switch to said storage manager and to said client and storage
35 devices.

1 29. The apparatus of claim 28 wherein each said port maintains its own redirection circuitry
2 and its own look up table storing redirection commands, and wherein said microprocessor in said

PATENT

3 storage manager server looks up the specific port to which a specific storage device involved in a
4 transaction is coupled and/or the specific port to which the client device involved in a specific
5 transaction is coupled and sends said one or more redirection commands for the transaction only
6 to the appropriate specific port.

1 30. The apparatus of claim 28 wherein each said port maintains its own look up table storing
2 redirection commands, and wherein said microprocessor in said storage manager server sends all
3 said redirection commands to all ports in said switch and wherein each port includes redirection
4 circuitry to store redirection commands in the port's look up table and use the data therein to
5 relabel for redirection data and transfer ready frames.

1 31. The apparatus of claim 28 wherein said switch maintains a shared look up table storing
2 all redirection commands received from said storage manager, said shared look up table for use
3 by all ports, and wherein said microprocessor in said storage manager sends all said redirection
4 commands to a port to which said storage manager server is coupled with instructions to forward
5 said redirection command for storage in said shared look up table, and wherein each said port of
6 said switch has circuitry to use said shared look up table to determine if an incoming frame or
7 packet needs to be relabelled for redirection and at least the ports coupled to said storage
8 manager include circuitry to store all said redirection commands in said shared look up table.

1 32. The apparatus of claim 28 wherein said microprocessor in said storage manager server is
2 programmed to sum the number of outstanding read requests that have not yet been fulfilled,
3 and only issues redirection commands when the number of outstanding read requests exceeds a
4 threshold which is either fixed or programmable and said size is greater than said redirection size
5 criteria.

1 33. The apparatus of claim 28 wherein said microprocessor in said storage manager server is
2 programmed to monitor the time between receipt of a read or write request and completion
3 thereof, and only issues redirection commands when the time to fulfillment exceeds a threshold
4 which is either fixed or programmable and the size of said transaction is greater than said
5 redirection size criteria.

1 34. A packet or frame switched network including one or more client devices and one or
2 more storage devices, each of which is coupled to a packet or frame switch, comprising:
3 a switch having conventional routing circuitry therein and further comprising:
4 shared redirection circuitry;

PATENT

5 a shared look up table;

6 one or more ports in said switch which are structured to operate to receive
7 relabelling and redirection and purge commands from a storage manager and to
8 forward them to said shared redirection circuitry for storage in said shared look up
9 table and structured to forward data and transfer ready frames to said shared
10 redirection circuitry and to forward all other types of packets or frames to said
11 conventional routing circuitry;

12 and wherein said shared redirection circuit includes circuitry to receive said
13 redirection commands and store them in said shared look up table and to receive
14 incoming data or transfer ready packets or frames from said ports and compare
15 the address data therein to the old address data in said look up table, and, if
16 there is a match to any of the old address data in said look up table functioning
17 to substitute the new address data from the matching entry in said look up table
18 for the old address data in the incoming packet or frame and pass the packet or
19 frame to said conventional routing circuitry, said shared redirection circuit also
20 structured to receive said purge commands and use the information therein to
21 purge from said shared look up table the old and new address data of redirection
22 commands for completed read or write transactions; and

23 a storage manager including conventional cache memory and circuitry programmed or
24 structured to receive read or write requests from said one or more client devices identifying
25 specific data to be read and to map the identified data to a storage device and, for each read or
26 write request, send a corresponding read or write requests to said storage device to get the
27 transaction started, and for storing the most frequently requested data blocks from read
28 transactions in said cache memory and programmed or structured to monitor read requests for
29 read requests for blocks of data that are stored in said cache memory, and, if a requested block
30 of data is stored in said cache memory, for sending the requested data from the cache memory
31 back to the client device which requested the data instead of sending a read request to a
32 storage device that stores the requested data, and further comprises circuitry to decide whether
33 to redirect a particular read or write request based upon fixed or programmable redirection criteria,
34 and programmed or structured to determine from the frame or packet header information of said
35 packet or frame the specific client device from which the request came, and, if said read or write
36 request is not to be redirected, for performing conventional processing to complete said read or
37 write request, but if said read or write request is to be redirected, for generating and sending one
38 or more redirection commands to said switch for storage in said shared look up table maintained
39 therein, said redirection command containing old address data and new address data that will
40 cause said shared redirection circuitry that receives a data or transfer ready packet or frame that

PATENT

4 1 is part of the transaction to be redirected to relabel said packet or frame with new address data
4 2 and forward said packet or frame to said conventional routing circuitry such that said packet or
4 3 frame gets routed so as to bypass said storage manager but so as to look like it came from said
4 4 storage manager, and programmed or structured to monitor status frames to determine when a
4 5 redirected transaction has been completed and sending purge commands to said switch to purge
4 6 the old and new address data from said shared look up table for read and write transactions that
4 7 have been completed; and
4 8 a data path coupling said switch to said client and storage devices and to said storage
4 9 manager.

1 35. A packet or frame switched network including one or more client devices and one or
2 more storage devices, comprising:
3 a data path coupled to said client devices and said storage devices;
4 a storage manager coupled to said data path;
5 a switch having conventional routing circuitry and including:
6 one or more ports in said switch which contain redirection circuitry to receive
7 relabelling and redirection commands from a storage manager and to store the redirection
8 address data containing old address data and new address data contained therein in at
9 least one look up table and to receive incoming data and transfer ready packets or
10 frames and comparing the address data therein to the old address data in said look up
11 table and to compare packet or frame type information included in a header of said
12 packet or frame, and, if there is a match to any of the old address data in said look up
13 table and the type information matches the types of packets or frames that are to be
14 redirected, functioning to substitute at least the new address data from said look up table
15 from the entry which includes old address data that matched the address data in the
16 incoming packet or frame and pass the packet or frame to said conventional routing
17 circuitry and performing conventional processing to route all other packets or frames to
18 said conventional routing circuitry; and
19 wherein said storage manger includes:
20 means for receiving write and read requests from any of said one or more client devices
21 identifying specific data to be read or written and for mapping said read and write requests to
22 particular storage devices, and for sending corresponding read or write requests to said
23 storage devices to get the transaction started, and for determining if the transaction is to be
24 redirected, and, if not, for performing conventional processing to complete the transaction,
25 but if said transaction is to be redirected, for analyzing said read and write requests to
26 determine how many blocks of data are to be read or written, and, if the number of blocks is

PATENT

greater than a threshold, which can be fixed or programmable, for mapping the requested blocks to a storage device which stores the requested data and for generating and sending one or more redirection commands to said switch for storage in said at least one look up table maintained therein, said redirection commands containing old address data and new address data that will cause the redirection circuitry in the port that receives a data or transfer ready packet or frame to relabel said packet or frame with new address data that will cause said packet or frame to be routed by said conventional routing circuitry so as to bypass said storage manager but look like it came from said storage manager, and for monitoring status packets or frames to determine when a redirected transaction has been completed and sending purge commands to said switch to purge the old and new address data from said at least one look up table for read and write transactions that have been completed.

36. A process of redirecting data frames or packets in a packet or frame switched network comprising the steps:

- receiving at a storage manager from a client computer one or more read requests;
- receiving from a storage manager at a packet or frame switch at least one redirection command to redirect the data frames or packets of at least some of said read requests originated by a client device and storing said redirection commands;
- transmitting corresponding read requests from said storage manager to one or more storage devices;
- responding to said read requests by transmitting frames or packet containing the requested data addressed to said storage manager; and
- in said switch, intercepting said frames or packets of data responding to read requests, and, if a redirection command is stored for a read transaction which certain packets or frames containing the requested data are a part of, changing the header of the frames or packets of said particular read transaction which contain the requested read data so as to reroute them through said switch so as to bypass said storage manager but relabeling the headers thereof such that the rerouted packets or frames appear to originate from said storage manager as a response to said read request.

37. A process of redirecting data frames or packets in a packet or frame switched network, said process carried out in a switch and comprising the steps:

- receiving from a storage manager at a packet or frame switch redirection commands to redirect the data frames or packets which are responses to at least some of said read requests originated by a client device and storing said redirection commands; and

PATENT

6 in said switch, intercepting frames or packets of read data responding to said at least
7 some read requests, and, if a redirection command is stored for a read transaction which
8 certain packets or frames containing the requested read data are a part of, changing the
9 header of the frames or packets of said particular read transaction which contain the
10 requested read data so as to reroute them through said switch so as to bypass said storage
11 manager but relabeling the headers thereof such that the rerouted packets or frames appear
12 to originate from said storage manager as a response to said read request.

1 38. A process for controlling redirection of data frames or packets in a packet or frame
2 switched network, said process carried out in a storage manager and comprising the steps:
3 receiving at a storage manager from a client computer one or more read requests;
4 composing and transmitting to a packet or frame switch redirection commands to redirect
5 the data frames or packets of at least some of said read requests originated by a client
6 device such that said switch relabels data frames for routing so as to bypass said storage
7 manager in travelling to the client device which made the read request but make said
8 relabelled data frame look like it came from said storage manager;
9 transmitting said read requests from said storage manager to one or more storage
10 devices.

1 39. A process of redirecting data frames or packets in a packet or frame switched network,
2 said process carried out in a switch and comprising the steps:
3 receiving from a storage manager at a packet or frame switch redirection commands to
4 redirect the data frames or packets which are responses to at least some of said read
5 requests originated by a client device and storing said redirection commands; and
6 in said switch, intercepting frames or packets of read data responding to said at least
7 some read requests, and, if a redirection command is stored for a read transaction which
8 certain packets or frames containing the requested read data are a part of, changing the
9 header of the frames or packets of said particular read transaction which contain the
10 requested read data so as to reroute them through said switch so as to bypass said storage
11 manager but relabeling the headers thereof such that the rerouted packets or frames have
12 the Port_ID of the client that made the read request as the destination address and the
13 Port_ID of the storage device as the source address and have the originator exchange ID
14 originally assigned by said client; and
15 in said client, receiving said frames and mapping them to the original read requests sent
16 to said storage manager.

PATENT

1 40. A process of redirecting data frames or packets in a packet or frame switched network
2 comprising the steps:
3 receiving at a storage manager from a client computer a write request;
4 receiving from a storage manager at a packet or frame switch at least two redirection
5 commands to redirect the data and transfer ready frames or packets of said write request and
6 storing said redirection commands;
7 transmitting a corresponding write request from said storage manager to a storage device
8 to get the transaction started;
9 at said storage device, responding to said write request by transmitting one or more
10 transfer ready frames from said storage device when it is ready to receive at least some of the
11 write data, said transfer ready frame(s) or packet(s) addressed to said storage manager; and
12 in said switch, intercepting said transfer ready frame(s) or packet(s) of data responding to
13 write request, and, if a redirection command is stored for a write transaction which said
14 transfer ready frame(s) or packet(s) are a part of, changing the header of the transfer ready
15 frame(s) or packet(s) of said write transaction to be redirected and routing transfer ready
16 frame(s) or packet(s) so as to reroute them through said switch so as to bypass said storage
17 manager but so as to appear that transfer ready frame(s) or packet(s) originated from said
18 storage manager;
19 at said client device, receiving transfer ready frame(s) or packet(s) and responding by
20 sending one or more frames or packets of write data to said switch;
21 at said switch, receiving said one or more frames or packets of write data and if a
22 redirection command is stored for a write transaction which said write data frame(s) or
23 packet(s) are a part of, changing the header of the write data frame(s) or packet(s) of said
24 write transaction to be redirected and routing said write data frame(s) or packet(s) so as to
25 bypass said storage manager but so as to appear that write data frame(s) or packet(s)
26 originated from said storage manager; and
27 continuing the process re-routing said transfer ready and write data frames until the entire
28 write transaction is completed.

1 41. A process of redirecting data frames or packets in a packet or frame switched network,
2 said process carried out in a switch and comprising the steps:
3 receiving from a storage manager at a packet or frame switch one or more redirection
4 commands to redirect the data and transfer ready frames or packets which are part of a write
5 transaction originated by a client device and storing said redirection commands; and
6 in said switch, intercepting write data and transfer ready frames or packets which are part
7 of a write transaction to be redirected and, if one or more redirection commands have been

PATENT

8 stored for the write transaction of which said write data and transfer ready frames or packets
9 are a part, changing the header of the write data and transfer ready frames or packets so as
10 to reroute them through said switch so as to bypass said storage manager but so as to look
11 like they came from said storage manager and so as to have the originator exchange ID
12 originally assigned by said client to the write transaction; and
13 transmitting said write data and transfer ready frames or packets to said client or said
14 storage device as appropriate.

1 42. A system comprising:

2 a data path or paths;
3 one or more client devices coupled to said data path(s);
4 one or more storage devices coupled to said data path(s);
5 one or more switch means coupled to said data path(s) for receiving packets or frames
6 and routing them appropriately according to header address information, and for receiving
7 redirection commands and storing them in one or more lookup tables and for using address
8 data in said redirection commands to redirect predetermined packets or frames, and for
9 receiving purge commands and for using the data in said purge command to purge from
10 said one or more lookup tables said redirection commands for read and/or write transactions
11 that have been completed ; and

12 storage manager means for receiving read and/or write requests from said one or more
13 client devices and for determining by any means whether or not to redirect said read and/or
14 write requests, and, if a read and/or write request is to be redirected, for generating and
15 sending to said switch one or more redirection commands to cause predetermined packets or
16 frames of said read and/or write transaction to be redirected, and for determining when a
17 read and/or write transaction that has been redirected has been completed and for
18 generating one or more purge commands pertaining to said read and/or write transactions
19 that have been completed, and sending said purge commands to said switch to cause it to
20 purge the redirection commands of completed transactions, and, if a read or write transaction
21 is not to be redirected, for processing it conventionally.

1 43. A switch for a packet switched or frame switched network, said switch having
2 conventional routing circuitry, said network having client devices and storage devices and a
3 storage manager that decides to redirect read and/or write transactions on any basis so as to
4 virtualize said storage devices, said switch comprising:

5 one or more ports which are structured to operate to receive one or more redirection
6 commands from said storage manager or another switch in said network coupled to said

PATENT

7 storage manager containing old address data and new address data, said ports containing
8 circuitry to store said old address data and new address data of said one or more redirection
9 commands in one or more look up tables, said switch containing circuitry functioning to
10 compare packet or frame type information included in a header of at least some packets or
11 frames arriving at a port coupled to a client device or a storage device to determine if the
12 packet or frame is of a type which is to be relabelled and redirected, and, if the packet or
13 frame is of a type to be relabelled and redirected, comparing at least some of the old address
14 data therein to the old address data entries for one or more redirection commands stored in
15 one or more of said look up tables, and if there is a match to any of the old address data of
16 said one or more redirection commands, relabelling said packet or frame by substituting the
17 corresponding new address data from said matching redirection command(s) for said old
18 address data in said packet or frame header, said new address data being such as to cause
19 said relabelled packet or frame to be routed so as to bypass said storage manager but to
20 look like it came from said storage manager, and then passing said relabelled packet or frame
21 to said conventional routing circuitry for forwarding to the destination identified in said new
22 address data, each port further comprising circuitry to recognize incoming status frames
23 indicating that a particular read or write transaction has been completed and to respond
24 thereto by automatically purging from one or more lookup tables one or more redirection
25 commands that pertain to the completed transaction of which said status frame is a part and
26 route said status frame to a port of said switch coupled to a storage manager, and wherein
27 the sequence in which the type of incoming packet is determined and the old address data
28 lookup is performed is not critical.